N89-20062

COMPUTER TECHNOLOGIES AND INSTITUTIONAL MEMORY

Final Report

NASA/ASEE Summer Faculty Fellowship Program-1988

Johnson Space Center

Prepared By:	Christopher Bell
Academic Rank:	Assistant Professor
University & Department:	Illinois State University Department: Applied Computer Science Normal, Illinois 61761-6901
Prepared By:	Roy Lachman, Ph.D.
Academic Rank:	Professor of Psychology
University & Department:	University of Houston Department: Psychology Houston, Texas 77004

NASA/JSC

Directorate:	Space and Life Sciences
Division:	Man-Systems Division
Branch:	Crew Station Branch
JSC Colleague:	Barbara Woolford, MA Frances E. Mount, MS
Date Submitted:	August 12, 1988
Contract Number:	NGT 44-005-803

ABSTRACT

NASA programs for manned space flight are in their 27th year. Scientists and engineers who worked continuously on the development of aerospace technology during that period are approaching retirement. The resulting loss to the organization will be considerable. Although this problem is general to the NASA community, the problem was explored in terms of the institutional memory and technical expertise of a single individual in the Man-Systems Division. The main domain of the expert was spacecraft lighting, which became the subject area for analysis in these studies. The report starts with an analysis of the cumulative expertise and institutional memory of technical employees of organizations such as NASA. A set of solutions to this problem are examined and found inadequate. Two solutions were investigated at length: hypertext and expert systems. Illustrative examples were provided of hypertext and expert system representation of spacecraft lighting. These computer technologies can be used to ameliorate the problem of the loss of invaluable personnel.

DESCRIPTION OF THE PROBLEM*

Personnel represent a major share of the resources of any organization. As these individuals interact with others in and outside of their organization, as problems are confronted and solutions identified, the organization's personnel increase in value to the organization. The solutions, the interactions, the actions of the personnel of the organization during their association with it contribute to the institutional memory of the organization. Each employee is privy to part of this memory to the extent the employee has participated in the organization's activities. Much of this institutional memory resides only within the employees. As such it is at risk with the employee. When any employee resigns, retires or dies, their unique part of the institutional memory is lost.

This paper examines this problem as it confronts NASA. It describes the early stages of an effort - using the tools of expert system development and hypertext data management schemes - at NASA-Johnson Space Center/Life Sciences Directorate to capture some of this institutional resource before it is lost. While NASA is a unique organization in many respects, its efforts to confront and resolve this problem may provide some guidance for other organizations which are encountering similar problems.

Loss of Institutional Memory

Institutional memory is the history of the activities and interactions which have occurred within an organization. Some of this history is recorded in the form of memos, papers, research studies and other archival material, but a large share is resident in the ephemeral memories of employees of the organization. Much of the information is common to more than one individual, but each individual has a unique combination of elements of the larger whole. They can also claim some unique elements of the institutional memory. As these individuals are separated from the organization - through transfers, retirements or death - the unique components of the institutional memory which they possess are lost to the organization.

Not all individuals within the organization can be considered to possess equally valuable components of the institutional memory. The corporate technical experts

*The order of authorship is alphabetic: C. Bell, specialized in the hypertext solution, R. Lachman in expert systems.

will be more highly valued than the corporate chauffeurs. Ideally, a process for the identification of the critically important experts and the capture of their critical experiences can be in place early in the expert's association with the organization. Such a process would greatly ease the problems of retaining within the organization the expert's uniquely valuable knowledge. The failure to identify these individuals prior to their leaving the organization will ensure loss of their unique components of the institutional memory.

NASA has a workforce which is rapidly approaching retirement age. Due to the unfortunate conjunction of similar hiring dates - associated with the start of the space program - and the number of individuals rolling over their military retirement into the federal system, some divisions (e.g.; Man-Systems Division) have over 50% of their technical and managerial personnel eligible for retirement. This is an aging population as well, subject to the impact of degenerative disease processes and common mortality risks.

Senior experts serve other roles in organizations beside those related to the application of their expertise and as a repository of institutional memory. They provide mentor services to junior personnel, acquainting new hires with successful and unsuccessful experimental and behavioral procedures. With the loss of the senior individuals the efficiency of the organization declines as non-productive techniques are retried, as mistakes are repeated by the less experienced new personnel.

Approaches to Solution of Problem

Several techniques have been tried for managing the problems associated with the loss of senior personnel. These can be grouped into those designed to reduce the rate of personnel loss and those designed to compensate for this loss. The organization can attempt to reduce the number of personnel lost to retirement by providing compensation - increased pay, more recognition, etc. - to these individuals. Appeals to corporate loyalty or patriotism may be tried to retain vital personnel. Legal adjustments - such as removal of mandatory retirement laws - have reduced arbitrary barriers. If retirement is not preventable, efforts can be made to reduce the impact of the expert's leaving. These efforts might take the form of long term post-retirement consulting assignments. A phased retirement package made of gradually extended vacation time allotments during the employee's last few years of service might help to reduce the impact to the organization of

sudden retirement.

Several approaches are available to compensate for the loss of personnel. To try to retain some elements of expertise at risk, an organization may assign apprentices to senior experts. These apprentices can reduce the impact of the loss of the expert by providing a partial replacement during a transitional period. New archival technologies are being introduced which will provide computer assistance to the problem. The organization can manage the problem in the traditional fashion of accepting the loss of the expert and their unique institutional memory components.

None of the retirement-rate reduction programs described compensate for the loss of experts due to age related or accidental processes. Consulting programs or phased retirement programs are weakened by two conflicting issues. To the extent that the expert is out of the normal flow of the organization's information and duties - referred to as the "loop" - the value of the expert's unique contribution is diminished. The more attempts are made to keep the expert in the "loop" the less others in the organization are able to accept and begin to compensate for the loss of the expert. There is a morale problem associated with excessively delayed retirement programs as well. The senior expert should be able to enjoy a healthy retirement period. Junior employees may begin to resent the fact that advancement routes are being blocked by the continued presence of the senior experts.

The use of apprentices, trainees or proteges would seem to give a semblance of continuity to activities requiring high levels of expertise. The success of this approach is dependent upon several factors. It assumes that there are interested trainees available. It assumes that the expert is sufficiently introspective to identify the critical features of themselves that make them valuable to the organization and is able to determine effective techniques to transfer these features to another. The expert needs to be able to verbalize his expertise and identify situations where it may be safely practiced by the trainee. These are not common skills nor are they easily learned. Finally it is unlikely that the expert can pass along the authority and knowledge associated with many years experience in a short time period.

Simple acceptance of continuing losses is the most commonly chosen alternative. Organizations such as NASA will find the cost of these losses very onerous. This leaves NASA with a strong motivation to explore

alternate archive generating and organizing procedures or other technological fixes.

An active archival program with a professional staff to gather, cross-reference and maintain significant documents is an expensive option. A difficulty in using such an approach for the problem of the loss of institutional memory through personnel attrition is that the material which is retained is the material which is retainable - printed text - and not the critical material of attempts successful or not, of judgements validated or disabused. These are retained by the expert and make up a crucial portion of the expert's unique component of the institutional memory. While these experiential items may be vital to the expert and their handling of new problems, they are not nearly as likely to appear in archival storage as memos or minutes of informational meetings. Had the expert maintained an active report or diary of problem solutions and attempted strategies, of backgrounds to decisions made and processes observed, the problems associated with both the archival and the apprenticeship approaches would be reduced. If the retiring expert can be induced to create a set of documents detailing the significant issues used to identify and solve problems, the thought processes and attack strategies needed, then the problems outlined in the preceding paragraphs will recede. It is unfortunate that most experts find it extremely difficult to accurately verbalize the techniques they use to identify solutions. Most are not even conscious of the procedures they are applying to the solution. Memory that may be critical to the decision making process may not be consciously available until needed for the solution of a particular problem. The reminiscence approach has a further negative associated with it. If a particular domain or pattern of problems has not been addressed recently, the information may not be brought to conscious memory in time to be added to the verbal record.

Even with complete memoirs of retirees, problems of organization, selection, access and structure remain. Access procedures must be simple enough to facilitate their use by remaining employees. They need to be robust enough to survive and be able to produce useful information for users who may not have the time to become facile in using a sophisticated data search package. Cross references between memoirs and other archival material need to be maintained, and a data dictionary of terms and synonyms established. The standard to work toward is an on-line system available to users within their desk environment or through simple telephone connections - equivalent in ease to contacting the former employee. It can be anticipated that a significant

portion of the material provided will be redundant, but a redundancy of information is to be preferred to a shortage of information. A winnowing process can control for information duplicated by other retirees or available from other sources.

Experts serve as more than information repositories. They achieve their status through the appropriate application of their knowledge to the situations they confront - they make decisions. While present software and theoretical activities point toward efforts to duplicate the broad decision making processes of experts, the most successful applications thus far have been those where the domain of interest is very narrow. Expert system tools have been most successfully applied in highly bounded situations, where the range of variables of interest can be anticipated and their interactions predicted and encoded. Unbound situations provide the greatest challenge to the extension of this technology. The problem of the loss of institutional memory and consultative resources is a very unbounded situation.

The replication of the tasks of the expert may be approached as a problem in artificial intelligence (i.e. as an expert system problem), or as a data base/archival problem. Given the specific needs of the person seeking to use the expert as a consultant, a traditional query langüage/key word - synonym archival approach will require substantial knowledge of the information domain prior to effective use. Many of the individuals contacting an expert express great unwillingness to become more familiar with the knowledge domain of the expert than they already were. Too often keyword searches result in too many "hits" to be useful to a simple informational search or result in too few "hits" because of a lack of overlap between the user's vocabulary and that of the expert. Scrolling is a very inefficient way to search a data base.

Given the difficulties outlined above for using standard archival approaches, the need for alternative approaches seems evident. We will examine two: hypertext and expert systems.

HYPERTEXT

If sophisticated data storage systems could be combined with easy to use linking or searching schemas, many of the difficulties of recovering information from a text base could be reduced. Such a combination of linked knowledge stores and an effective front end interface is available in present software shells under the generic

term hypertext or hypermedia.

> The concept of hyper-text is quite simple: Windows on the screen are associated with objects in a (knowledge) base, and links are provided between these objects. (Conklin, 1987).

General Description

The term hypertext is used in this paper to describe a set of techniques for creating and linking text-oriented nodes of information which may be accessed dynamically during any given trial. Each node represents an item of interest in the hypertext knowledge base. Links are established to other nodes containing information associated with the node. Links are conceptual pathways that are used to move from one node to another. A user browses in the hypertext knowledge base seeking information. This description of hypertext is meant to be general. Each experimental or commercial shell or developed package will differ depending upon the creator's personal view of the relative importance of various features. For an excellent review of the specific implementations of hypertext systems, see Conklin (1987).

Jumps between related nodes occur at the user's discretion. The user has the option of pursuing any line of inquiry by following links to nodes that may contain relevant information. While the links need to be pre-established, the exact path chosen through the hypertext knowledge base is at the user's discretion and thus not rigid. Links may be followed in both directions. Any given node may serve as the target for one or more links and may have no, one or several links leading from it. Frequently, a graph or map of the hypertext knowledge base is maintained by the system to serve the user as a graphic memory aid. Whether the underlying linkage structure should be hierarchical or not seems to be at the developers' discretion. Conklin (1987) presents arguments for both approaches. Although most users inherently favor a hierarchical structure some need for the ability to jump between structures seems to be supported. Users wishing to jump from one node to another may do so with minimal keystrokes. The software keeps track of the user's location and efficiently permits the user to quickly leave a node to move to a higher level (assuming a hierarchical structure) or to another node with ease. The simplest analogy to a large non- structured hypertext system is to visualize interlaced cobwebs with nodes of information at each of the cross points. The user is like the spider running along the web, each strand representing a link. This is in

contrast to traditional data base schemas where the pattern of linkage and of searching is rigidly structured (as in hierarchical and network designs) or designed for pattern matching (as in relational data base organizations).

Extraction of nodes

The creation of the hypertext oriented system requires the efficient extraction of nodes from the background material. Nodes are edited to represent the information they contain in simple useful text. Text quantity may correspond to a sentence, a paragraph or a screen. Given the desire to produce useful and frequently used reference materials nodes should be edited with a bias favoring the simple and the short over the long and convoluted. As the nodes are extracted, they are linked to related nodes. These links are established to provide possible routes of access for users. While it is impossible to anticipate all possible routes of user interest, it is suggested that the original source of the material might serve very nicely as an initial approximation of the pattern of relationships. Cross links are then made to relevant topics in nodes extracted from other documents. A cognitively manageable structure needs to be retained. Commonly, clues to links to other nodes have been in the form of embedded references within text or menu lists (see Fig. 1 and Fig. 2). Either the first line of the node is given or a node title is provided to aid the user in deciding whether to use a provided link to the node. Both approaches have support in commercial systems.

THE HYPERTEXT PROJECT

The role of the expert is both as a consultant and as a decision- maker. As a consultant the expert is able to bring together information from previous work which may be relevant to the project at hand. The decision-making aspect involves using this information to reach a conclusion about the problem. The information management aspect of the expert can be approximated using hyper-text oriented hypertext knowledge base procedures. Information from the expert's descriptions of previous projects can be merged with data from other sources to approximate the knowledge of the expert, and then linked to provide a user with access to the information of the expert. Suggested sources of material to put into the hypertext knowledge base would include retirees' observations, technical reports and other collateral material. This hypertext knowledge base would be tied to a user

ORIGINAL PAGE IS
OF POOR QUALITY

```
                                                  LOCAL      8 OUTPUTS
                                            CENTRAL OFFICE
                                             DEPARTMENTS
                                             FACULTY
                                             GRADUATE ASSISTANTS
                                             CIVIL SERVICE
                                             NEW FACULTY
                                             CAMPUS LOCATION
                                             PHONE NUMBERS
CAST FACULTY HYPERTEXT

LINK: Main       Edit      File      Option      Prioritize       X
      Random      Sort      Tag      Untag       View      WordRef
Returns to the Main-menu
TEXT      CAPS 1     460 NODES   1708 LINKS      354K Bytes A:\CASTDIR3
```

Figure 1. - Menu oriented hypertext screen.

```
        to users the actual methods that will lead to a particular decision.
        With hypertext, users directly participate in each and every decision
        that leads to a particular expertise.  This process of openly
        displaying the structure <FILE75 STRUCTURE> and uses of knowledge
        readily transmits it to users of hypertext systems.

SENSITIVITY ANALYSIS   Experts machines usually provide a single answer
        supported perhaps by a confidence factor such as 82 percent certainty
        (whatever that means).  Hypertext systems allow users to rapidly test
        alternative paths <FILE56 BROWSE> to see how sensitive the advice may
        be to changes in the initial assumptions.

In my opinion, given the years of AI promises and the relative lack of
success in creation and delivery of workable experts systems, the most
important pocketbooks are rapidly closing against further use of computer
processing of rule-based systems as a method of vending expertise from a
disk.

Generally, over the last 30 years, technologists haven't made operations
research practical.  Consequently, if you can't build machines that are
effective in quantitative reasoning, how can you build machines that are
effective in subjective reasoning.  Making subjective decision machines work
is at least several magnitudes more difficult.  <FILE27 REASONS>

-- More --   Page 4  Columns 0 - 79   Rmargin: 77   FILENAME: file21
```

Figure 2. - Embedded link hypertext screen.

interface designed for easy, effective interactions. With the passage of time the organization will progress beyond the areas of the technical competency of the retired expert. The number of archival inquiries can be expected to drop. The nature of those inquiries can be expected to change from specific information seeking to general viewpoint or attack strategy issues. This represents the best presently available solution to the problems associated with the need for a partial replacement for retiring experts.

The expert's interpretation of the information cannot be duplicated, nor can the expert's judgement as to the relative importance to attach to the various nodes. The user still must make a judgement as to the weight to put on the information so presented. The advantage to the user is that the information is available quickly. Associated information is rapidly accessed and relative importance to the project at hand can be judged. The user is integrated with the computer hypertext knowledge base, rapidly pursuing leads and identifying interesting nodes and rejecting those of little interest. While the expert is not available, the user is able to make approximation of the expert by accessing the same information, linked in the same manner as the expert would have linked it.

NASA has a critical need to establish protocols for capturing the expertise of retiring expert personnel. To this end an individual was identified who has served as a consultant in a variety of areas including lighting and window design. A brief sampling of the telephone calls received by this individual supported the claim of being consulted over fifty times per month on issues related to areas of expertise. This particular individual is also very verbal and it was hoped that this would also help, given the prototype nature of the project.

Procedure for Gathering Information

Initial data was gathered during interviews with the subject. The interviewers quickly discovered that it was necessary to isolate the subject from his normal work environment as the demands on the subject for consultation were frequent. Interviews were transcribed and the transcription files served as the core for the hypertext knowledge base.

Initial interviews were minimally directed. The subject was asked to discuss his experiences with lighting issues during the development of various manned space efforts. Review of the interviews identified lacunae which were used as prompting points in

subsequent interview sessions. There were two objectives being met during the interview series. One was to provide a core of material for the development of the hypertext knowledge base. The interviews also served as a means to isolate some examples of decision making that could serve as models during the development of the proto-type expert system. Both interviewers were present during the sessions. Some conflicts developed as to directions for follow-up questions or deciding the point at which a narrative needed to be redirected. Sufficient information was gathered for both efforts.

Development of Hypertext Knowledge Base

Given the goal of creating a useful and used product, the hypertext knowledge base uses the narrative material as one of several important sources. The second major source is a NASA document titled Man-Systems Integration Standard. This document is a product of the Man-Systems Division and serves as a summary of human factors related information for NASA. Other sources useful to this project are published reference material dealing with vision and light. Material from the latter sources can be updated as new interpretations of data cause changes to be made in the scientific publications relevant to the topic area.

The hypertext document being produced for NASA uses the commercial hypertext software shell HOUDINI. Efforts are being directed toward examining issues related to determining the most effective interface for NASA personnel and other technically oriented users.

Questions to be addressed include:

How much freedom of movement - cross-linkages as opposed to hierarchical links - is needed by the casual user in contrast to the experienced user?

Should the user be able to annotate the hypertext knowledge base?

Should the user be able to add new links or nodes?

How does training effect user format preferences?

These will be discussed in future papers by the authors.

Users of hypertext systems have noted some operational problems. There is a significant cognitive overhead involved in keeping track of the information sought and in identifying nodes of immediately useful information giving a large number of irrelevant or partially relevant nodes. User disorientation - lost in hyperspace -

occurs when the user loses their cognitive map of the relationships between nodes. There is the risk of a combinatorial explosion of links if nodes are thoroughly cross-linked. Each of these areas will be actively addressed. Hypertext is not the final solution to the problems associated with the loss of institutional memory. It does present a methodology for reducing the effect to the organization of the loss of institutional memory through mortality of corporate experts. This is an expensive loss. Whether modern organizations can acknowledge the threat of this loss and make the required investment of time and effort to begin to compensate for it, is another question.

EXPERT SYSTEM APPROACHES

Project Scope and the Expert System Life Cycle

In the development of expert systems, the most important judgement is the one that determines the suitability of a task for expert system support. A positive decision in this regard represents the first in a series of sequential steps that are recommended for the development of an expert system prototype prior to the construction of a full-scale knowledge-based system (Waterman, 1986). Prototype development is one phase of the expert system life cycle. The life cycle consists of five major phases: system conceptualization and formulation of requirements, selection of software tools, developing the prototype, constructing the end product, and maintaining the product.

The summer project, according to its initial conception, revolved around the professional activities of the Man-Systems Division internal specialist in spacecraft lighting and windows, CW. Since CW is approaching retirement, it appeared prudent to examine the options available to NASA in any effort designed to preserve the operational experience he has accumulated during a period of about 30 years service to the organization. The expert system approach was among the two options that were selected for detailed examination in consultation with NASA colleagues. The scope of the project was conceived as more in the nature of a problem analysis and overview than the construction of an expert system prototype. Consequently, this part of the project is best characterized as a pre-prototype problem evaluation and the initial phase of prototype construction.

Typical knowledge-engineering methodology was used. The flexibility of the method is an asset but the concomitant absence of standards can be a serious impediment to the development of a knowledge-based system. Logically, the first step in the methodology consists of an evaluation of

the expertise of the domain expert selected for the project. It is essential to initially determine if the candidate consultant is in fact a genuine expert with respect to the subject domain. However, a necessary prerequisite for that judgement, as well as subsequent steps in prototype development, is an adequate level of familiarity with the domain of knowledge. Consequently several weeks were spent examining primary sources in lighting engineering and design (e.g. Kaufman and Christensen, 1984; 1987). The initial acquisition of domain vocabulary and conceptual structure was later evaluated, corrected, and extended by questions posed to CW during formal knowledge engineering sessions. During the period devoted to the study of the literature on illumination engineering, a log was kept of the consultations that the expert provided to other NASA professionals, the meetings that he attended, and the tests that he conducted.

The consultations are best characterized as problem solving sessions and covered a surprisingly broad range of topics. The expert devoted considerable time to analyzing competing designs for the cupola windows in the space station. In addition to forecasting general visual effects of window design modifications, he constructed AITOFF Equal Area Projection of the Sphere showing the kinds of visual function (foveal vision, color vision, binocular, etc) that could be lost as a result of several contemplated changes in window design. He also gave technical advice on attributes of space station windows such as size, number of panes, shape, replaceability, construction materials and selection criteria. The second area consisted of advice on space station vibroacoustics, including issues of subliminal vibration. Third, he developed TV overlays for viewing the grapler on the Canadian arm for handling payload during STS-24. Fourth, he conducted ambient lighting experiments that compared the relative efficacy of LCD and gas plasma displays for the on board laptop computers that display the location of the Orbiter. Fifth, he presented arguments dealing with the design of central lighting for workstation areas on the space station. (Someone wanted to change the overhead lighting specifications from 4800K to 3200K ostensibly to achieve superior resolution for TV monitoring of workstation activity. CW had photos in his files from previous tests that demonstrated the color temperature levels that produce better color resolution and rendering both to eye and camera.) Sixth, he collaborated on a study dealing with mercury containment in fluorescent lighting. The seventh and eight were consultations on color rendering in still photography and reflectivity of materials in the cargo bay of the Orbiter. The ninth problem that he worked on during the period of observation dealt with the design of an EVA helmet with LCD projection on its viewing surface.

The observations of consultative collaboration and advisory meetings led us to conclude that the expertise was indeed genuine. This conclusion was confirmed in interviews with a number of engineers and scientists at JSC who occasionally consult with CW. The extent of the expertise, however, was unexpected. Expert system technology does not work well with broad domains of knowledge and the issue of project suitability required that the project be narrowed to aspects of the domain of spacecraft lighting design. The domain represents one of CW's primary area of expertise. The design area may be segmented so that realizable goals may be set for pre-prototype problem evaluation and the initial steps of prototype construction.

System Conceptualization

The specification of system requirements and the description of user populations are part of the system conceptualization phase of prototype development. Several two hour long knowledge engineering sessions were conducted with the expert. These were followed by interviews with potential users of a lighting design expert system. The combination of literature review and interviews lead to the formulation of the first cut in the system design: IVA (Intravehicular Activity) and EVA (Extravehicular Activity) lighting design. Although the lighting principles may be the same for the two, the heuristics, rules, and tests required appear to be quite different. IVA was selected. The next cut was made in accordance with major areas of spacecraft activity: general circulation, habitation, and workstations. Emergency and portable lighting was added as a fourth category of lighting design. Habitation was subdivided into crew quarters, ward room, waste management, personal hygiene, and showers. Workstation was divided into exercisers, general maintenance, health care, galley, windows, and science labs. The development of a prototype is typically focused on a restricted subdomain so that the potential of a full scale expert system can be demonstrated. CRT and other display and control workstation components were not included in an overall menu design at this time because instrument panel lighting entails unique design principles that add considerable additional complexity. The problem was deferred for future analysis. Also, an enumeration of the full set of components of lighting design to be described in a prototype and later incorporated in a final system was not attempted.

The conceptualization phase of development also entails the identification of representational and problem-solving methodologies that are suitable for the task. These issues will be briefly discussed below.

The major conclusions of the system conceptualization and development analysis are:

1. The development of a Spacecraft Lighting Advisor is a major expert system construction task and accordingly requires the methodologies and effort that is typical in large-scale knowledge based system development. The design experience and advice appropriate for building small knowledge systems does not apply and can be seriously misleading for the task at hand.

2. The development of a full prototype system will require approximately one man-year provided that the one or more members of development staff have the necessary training or experience in artificial intelligence programming and in the psychology of knowledge engineering.

3. The finished system may take an additional 1.5 man-years or longer depending on the number of features included.

Tool Selection

The most expensive expert system development packages, in general, tend to be the most powerful and possess the widest array of representational technologies for development. They also include procedures for effective execution, as well as the maintenance and modification of the knowledge base. Unfortunately, the most expensive packages also require the longest period to master, and full mastery only occurs when the tool is used in a real world project. There is also the problem of matching the package to the problem. For example, Clancey (1985) maintains that "when presented with a given 'knowledge engineering tool' such as EMYCIN (van Melle, 1979), we are still hard-pressed to say what kinds of problems it can be used to solve."

Typically, a large-scale problem, such as the one under consideration, will require hybrid knowledge representation including rules, frames and other object-oriented devices. At this stage of the evolution of the technology, however, no one can definitively select a knowledge engineering tool as appropriate for a given type of problem, let alone the one that is optimum for that problem. Yet, effective tool selection remains a process that can place boundaries around a project and render it computationally and financially tractable.

Early in the project, a decision was made against allocating any of the ten week project period to an in-depth evaluation of tools, as the required analysis can consume a project of short duration leaving little time for anything else. Consequently, an expert system shell developed for an

advanced graduate course in knowledge-based systems (Lachman, 1989) was modified for the task. The first menu of the shell contains a selection that jumps the program to a module patched into the main program. The patched module was used to demonstrate a partial prototype for spacecraft lighting. Fig. 3 shows a menu nested at the second level and contained in the patched code. The knowledge base developed for the pre-prototype system can be readily ported to a commercial heuristic programming environment once a tool has been selected that was designed for building large-scale systems. The final programming environment will require demonstrable advantages for maintenance, explanation facility, control of reasoning, and knowledge base representation of the final product.

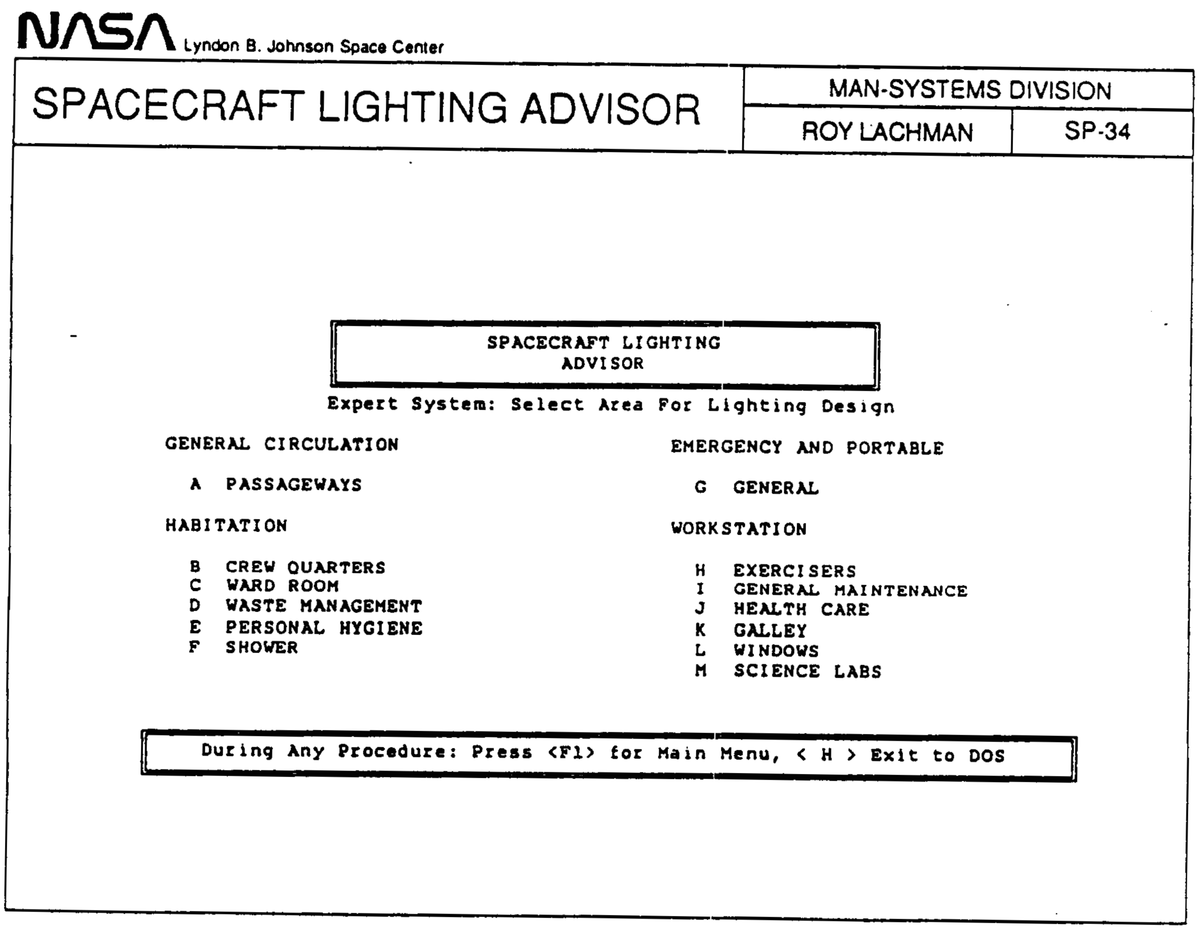

Figure 3.- A menu from the pre-prototype system.

The pre-prototype software includes elements of a backward chaining module that, when completed, would evaluate the design problem and interactively determine if the problem was appropriate for an expert system lighting advisor. This is not a trivial part of the project since design engineering can involve creative processes that cannot be implemented in automata at this time.

The overall plan for the expert system design is to try to duplicate processes generally used by human designers. The general theoretical orientation is the Newell and Simon (1972) information processing theory of problem solving. The problem space, in their theory, consists of a mental representation of the initial problem state, mental operators that change the problem state, and a termination criterion that is an acceptable solution to the problem at hand. The problem space consists of all the states or situations that, in principle, can be produced from the initial state by application of cognitive operators. A solution is any sequence of individual state-operator pairs that leads to a termination state. A problem-solving strategy is a method of search that identifies specific paths through the search space as worthy of consideration. Engineering design can be generally characterized in the search-space formulation. An expert system tries to capture from a human domain expert actual search-space strategies employed. A prototype should demonstrate the efficacy of the heuristic reduction of the search space and selective search operators that were copied from the human expert. It also should demonstrate the potential value of the final product to the sponsoring organization.

A pre-prototype system, in contrast, displays features in a pre-computational fashion that might be included in the final system. The goal is to demonstrate the logical adequacy of a given set of properties and a given approach. In addition, the pre-prototype may include a computational implementation of one or two of the features. Major properties of engineering design should be captured in the software. Engineers, typically, generate and test design alternatives for the search space of a given design specification. A rough design is produced and backtracking through the search space then produces incremental improvements in the design.

The simplest situation of problem solving in engineering design occurs when a previous design is implemented without modification; specifications are tested and a single unidirectional path through the search space is produced. Because of the project's time constraints, that type of solution was selected for the computational demonstration along with an uncomplicated area of lighting design on a spacecraft, the crew quarters. A sample

```
07-30-1988     16:16:19                    VER=S14 RULE-BASE=SLA2.RB

Try To establish one of the following hypotheses from the rule base:
 1 =NO-DESIGN-CHANGE:USE-STANDARD-15WATT-LUMINAIRE:TWO-PARALLEL-FIXTURES
 2 =ADAPTIVE-DESIGN:CHANGE-LOCATION-OF-FIXTURES
 3 =VARIANT-DESIGN:CHANGE-NUMBER-OF-FIXTURES-AND-LOCATION
 4 =AN-ORIGINAL-DESIGN-REQUIRED:REDESIGN-LUMINAIRE

ENTER HYPOTHESIS NO. --->1

The System is trying to prove goal NO-DESIGN-CHANGE:USE-STANDARD-15WATT-
LUMINAIRE:TWO-PARALLEL-FIXTURES
The System is trying to prove goal STANDARD-CREW-QUART.DIMENSIONS
The System is trying to prove goal STANDARD-D1
The System is trying to prove goal DIMENSION1(HEIGHT)=80IN.
All relevant rules have fired, can't prove DIMENSION1(HEIGHT)=80IN., which
is not in STM-DB nor the consequent of a Rule.

IS THIS TRUE: DIMENSION1(HEIGHT)=80IN.  (Y:N:D:W) --->Y

RULE  5 Deduces STANDARD-D1
Short-term memory (DATA BASE) now contains:
DIMENSION1(HEIGHT)=80IN. STANDARD-D1
The System is trying to prove goal STANDARD-D2
The System is trying to prove goal DIMENSION2(WIDTH)=41.5IN.
All relevant rules have fired, can't prove DIMENSION2(WIDTH)=41.5IN.

IS THIS TRUE: DIMENSION2(WIDTH)=41.5IN.  (Y:N:D:W) --->Y

RULE  6 Deduces STANDARD-D2
Short-term memory (DATA BASE) now contains:
DIMENSION1(HEIGHT)=80IN. STANDARD-D1 DIMENSION2(WIDTH)=41.5IN. STANDARD-D2

The System is trying to prove goal STANDARD-D3
The System is trying to prove goal DIMENSION3(DEPTH)=39.5IN.
All relevant rules have fired, can't prove DIMENSION3(DEPTH)=39.5IN.

IS THIS TRUE: DIMENSION3(DEPTH)=39.5IN.  (Y:N:D:W) --->Y
```

Figure 4. - A sample run of the pre-prototype expert system.

IS THIS TRUE: STANDARD-WEIGHT-CONSTRAINTS-APPLY <Y:N:D:W> --->Y

The System is trying to prove goal STANDARD-HEAT-CONSTRAINTS-APPLY
All relevant rules have fired, can't prove STANDARD-HEAT-CONSTRAINTS-APPLY, which is not in STM-DB nor the consequent of a Rule.

IS THIS TRUE: STANDARD-HEAT-CONSTRAINTS-APPLY <Y:N:D:W> --->Y

The System is trying to prove goal STANDARD-POWER-CONSTRAINTS-APPLY
All relevant rules have fired, can't prove STANDARD-POWER-CONSTRAINTS-APPLY

IS THIS TRUE: STANDARD-POWER-CONSTRAINTS-APPLY <Y:N:D:W> --->Y

RULE 9 Deduces STANDARD-WEIGHT:HEAT:POWER-CONSTRAINTS
Short-term memory (DATA BASE) now contains:
DIMENSION1(HEIGHT)=80IN. STANDARD-D1 DIMENSION2(WIDTH)=41.5IN. STANDARD-D2 DIMENSION3(DEPTH)=39.5IN. STANDARD-D3 STANDARD-CREW-QUART.DIMENSIONS STANDARD-WEIGHT-CONSTRAINTS-APPLY STANDARD-HEAT-CONSTRAINTS-APPLY STANDARD-POWER-CONSTRAINTS-APPLY STANDARD-WEIGHT:HEAT:POWER-CONSTRAINTS
RULE 1 Deduces NO-DESIGN-CHANGE:USE-STANDARD-15WATT-LUMINAIRE:TWO-PARALLEL-FIXTURES
Short-term memory (DATA BASE) now contains:
DIMENSION1(HEIGHT)=80IN. STANDARD-D1 DIMENSION2(WIDTH)=41.5IN. STANDARD-D2 DIMENSION3(DEPTH)=39.5IN. STANDARD-D3 STANDARD-CREW-QUART.DIMENSIONS STANDARD-WEIGHT-CONSTRAINTS-APPLY STANDARD-HEAT-CONSTRAINTS-APPLY STANDARD-POWER-CONSTRAINTS-APPLY STANDARD-WEIGHT:HEAT:POWER-CONSTRAINTS NO-DESIGN-CHANGE:USE-STANDARD-15WATT-LUMINAIRE:TWO-PARALLEL-FIXTURES

FINAL RESULT: NO-DESIGN-CHANGE:USE-STANDARD-15WATT-LUMINAIRE:TWO-PARALLEL-FIXTURES

07-30-1988 16:16:51 NORMAL TERMINATION

RULE 7 Deduces STANDARD-D3
Short-term memory (DATA BASE) now contains:
DIMENSION1(HEIGHT)=80IN. STANDARD-D1 DIMENSION2(WIDTH)=41.5IN. STANDARD-D2 DIMENSION3(DEPTH)=39.5IN. STANDARD-D3
RULE 8 Deduces STANDARD-CREW-QUART.DIMENSIONS
Short-term memory (DATA BASE) now contains:
DIMENSION1(HEIGHT)=80IN. STANDARD-D1 DIMENSION2(WIDTH)=41.5IN. STANDARD-D2 DIMENSION3(DEPTH)=39.5IN. STANDARD-D3 STANDARD-CREW-QUART.DIMENSIONS
The System is trying to prove goal STANDARD-WEIGHT:HEAT:POWER-CONSTRAINTS
The System is trying to prove goal STANDARD-WEIGHT-CONSTRAINTS-APPLY
All relevant rules have fired, can't prove STANDARD-WEIGHT-CONSTRAINTS-APPLY, which is not in STM-DB nor the consequent of a Rule.

Figure 4.(continued) - A sample run of the pre-prototype expert system.

demonstration run is shown in Fig 4.

It should be possible to develop a spacecraft lighting advisor in an incremental fashion starting with the least problematic design areas and working up to the most difficult. Successful development of a heuristic system that evaluates the difficulty and feasibility of design problems for solution by an expert system is not guaranteed at this stage in the evolution of the technology.

SUMMARY

In this paper, we have examined the problems associated with the loss of specialized expertise and institutional memory through the attrition of personnel. The costs to the NASA organization were examined and several possible solutions were critiqued. Two solutions were selected for detailed examination: hypertext and expert systems. A paradigmatic case was selected in the person of an in-house expert on spacecraft lighting and windows. Sample computational implementations were produced and described. These were determined to offer reasonable solutions to at least part of the problem.

REFERENCES

Clancey, W. J. (1985). Heuristic classification. Artificial Intelligence, 27, 289-350.

Conklin, Jeff (1987). Hypertext: An Introduction and Survey. IEEE Computer, 17-40, (September).

Kaufman, J. E. and Christensen, J. F. (1984). (Eds.) IES Lighting Handbook: Reference Volume. New York: Illumination Engineering Society of North America.

Kaufman, J. E. and Christensen, J. F. (1987). (Eds.) IES Lighting Handbook: Application Volume. New York: Illumination Engineering Society of North America.

Lachman, R. (1989). Expert Systems: A cognitive science approach. Behavior Research Methods, Instrumentation, and Computers, 21, in review.

Newell, A. and Simon, H. A. (1972). Human problem solving. Englewood Cliffs, NJ: Prentice-Hall.

Van Melle, W. A domain-independent production system for consultation programs. Proceedings of the Sixth International Joint Conference on Artificial Intelligence. Pages 815-823, August, 1981.

Waterman, D.A. (1986) A Guide to Expert Systems. Reading, MA: Addison-Wesley.